

The Adaptive Enterprise and HP OpenVMS

A White Paper describing how HP OpenVMS is addressing rapidly changing business demands, now and in the future

October, 2004



Executive Summary	2
Overview	2
Introduction	2
Business Strategy and Processes	2
Adaptive Enterprise Design Principles.....	4
Adaptive Enterprise Components	6
Virtualization	6
Applications	6
OpenVMS Web Services	7
Infrastructure Services.....	7
Continuous and Secure Operations	8
Security	9
Performance	9
Clustering, High Availability and Manageability	9
OpenView.....	9
Dynamic Resource Optimization.....	10
Virtual Server Environment - OpenVMS Galaxy	10
Workload Manager	11
Global Workload Manager.....	11
Process Resource Manager – OpenVMS Class Scheduler.....	11
Utility Pricing	11
Instant Capacity On Demand.....	11
Pay Per Use	11
Per Processor Licensing	12
Foundation for the Future	12
The OpenVMS Adaptive Enterprise Timeline	13
Conclusion.....	15
For more information.....	15
Call to action	15

Executive Summary

This white paper outlines how HP OpenVMS contributes to HP's Adaptive Enterprise initiative, so that customers can begin to plan integration of their current and future OpenVMS environments into the Adaptive Enterprise.

Overview

The following topics are covered:

- An introduction to the Adaptive Enterprise vision and why HP has launched this initiative
- Adaptive Enterprise areas of focus
- Products and Services that turn the vision into tangible benefits
- How OpenVMS on HP AlphaServer systems already has solutions for many of these areas
- Commitments that OpenVMS has made, based on the HP Integrity server family, to deliver more Adaptive Enterprise components in the future
- OpenVMS Adaptive Enterprise product delivery Timeline - AlphaServer systems and Integrity server family

Introduction

HP launched the Adaptive Enterprise initiative to the World in May, 2003. This was in response to the growing realization by companies that existing IT infrastructures are too inflexible to adequately respond to the demands of rapidly changing business and market needs. Business decisions that would take the Board perhaps a few days to agree on could take months to implement, with the main impediment being the incompatible silos of IT infrastructure deployed around the classic paradigm of disparate applications, databases, servers and networks. Companies need to be agile in their response to rapidly changing business/legislative/global demands and this is what the Adaptive Enterprise addresses.

HP is leading the delivery of computing resources as services, based on a flexible, virtualized infrastructure. Deployed this way the response time taken by the IT infrastructure to changing business demands can be dramatically reduced, improving time to market and overall customer satisfaction and cost. It is imperative of course to continue to support current applications, maintain quality of service, control costs and manage risk.

Business Strategy and Processes

In the past, a company's business strategy was likely to be in a single direction with multiple contingency plans. The more far reaching the strategy the more critical to success was the investment and deployment of an IT infrastructure that was, typically, inflexible. With the Adaptive Enterprise, an IT strategy can be defined as a range of critical capabilities and assets, responding to multiple future business needs.

Business processes support the strategy and detail the methods of operation to realize the goals. Processes cover the full gamut of IT operations including all inputs, outputs - financial and operational - including resourcing and capacity planning.

In today's environment of budget restrictions and IT frugality, business strategies and processes exercise the IT infrastructure to the limit and beyond. This can be attributed, at least in part, to the rigidity of the infrastructure – e.g. sizeable compute resources in a given location that are dedicated

to running specific database applications, may only be utilized less than half the day. If these resources were able to be swiftly reconfigured and reallocated to support operations in another location many time zones away over night and then restored to their original state at the beginning of the next working day, the company would realise a far greater return on those resources and their investment in IT.

Research by HP and independent analysts has highlighted another outcome of the compartmentalization of compute resources. The vast majority of systems are under-utilised for most of their working life but are likely to be severely stretched at times of heavy loading, such as Month End or the Christmas Rush.

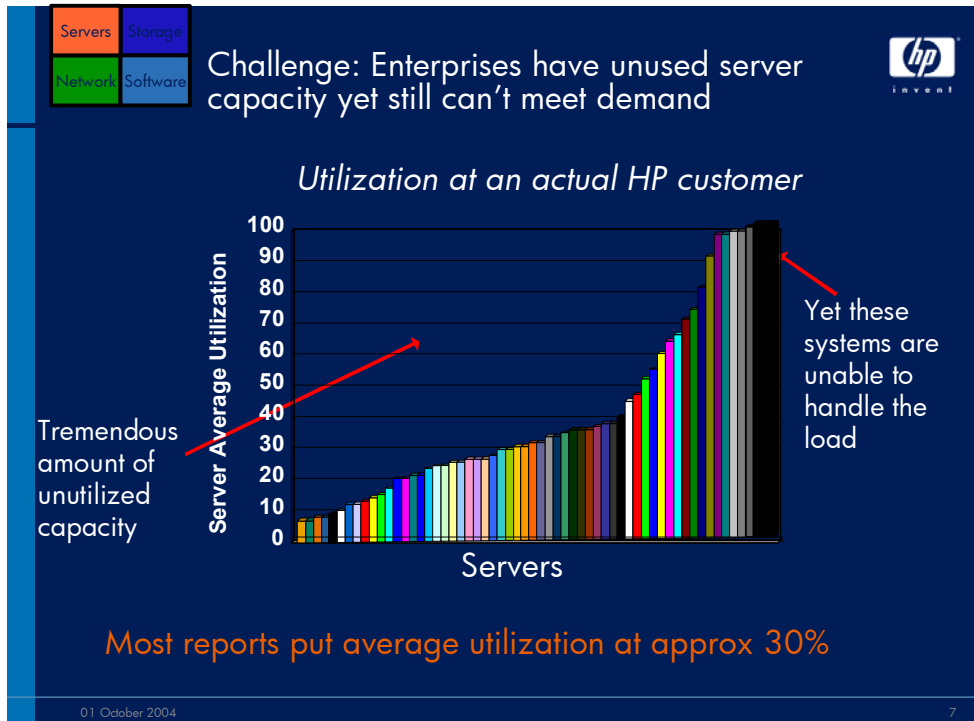


Figure 1 – Typical system utilization of a real world HP customer

This is what Gartner has to say about the current state of system utilisation:

Potential for Efficiency

"The potential for efficiency is quite impressive. On average, Intel servers are 15 percent loaded and RISC servers 40 percent or less. An average utilization rate of 80 percent — as for mainframes — is a long way off, but this landscape is actually extremely inviting to a consolidating and maturing industry"

from Utility Pricing Research Note – July 2004

To overcome this inefficiency and inflexibility HP have developed the Virtual Server Environment – VSE – that introduces the concept of pooling and sharing resources to cope with the fluctuating demands of myriads of work streams emanating from the various Business Units.

The Role of Virtualization in the Adaptive Enterprise



Business and IT synchronized to capitalize on change



Virtualization: an approach to IT that pools and shares resources so utilization is optimized and supply automatically meets demand

01 October 2004

4

Figure 2 – The role of Virtualization in the Adaptive Enterprise

Gartner also have a prediction about Virtualization that could seriously impact organizations that do not take notice of this emerging technology:

"Enterprises that do not leverage virtualization technologies will pay up to 40 percent more in acquisition costs by 2008, and roughly 20 percent more in administrative costs than enterprises that leverage virtualization technologies."

T. Bittman, Gartner Research—July 2003

Adaptive Enterprise Design Principles

Four principles underpin the Adaptive Enterprise:

- Simplification
- Standardization
- Modularity
- Integration

The four fundamental design principles are crucial to ensuring a tight link between IT and business. As these principles drive technology decisions, HP remains focused on delivering realizable and practical value for our enterprise customers.

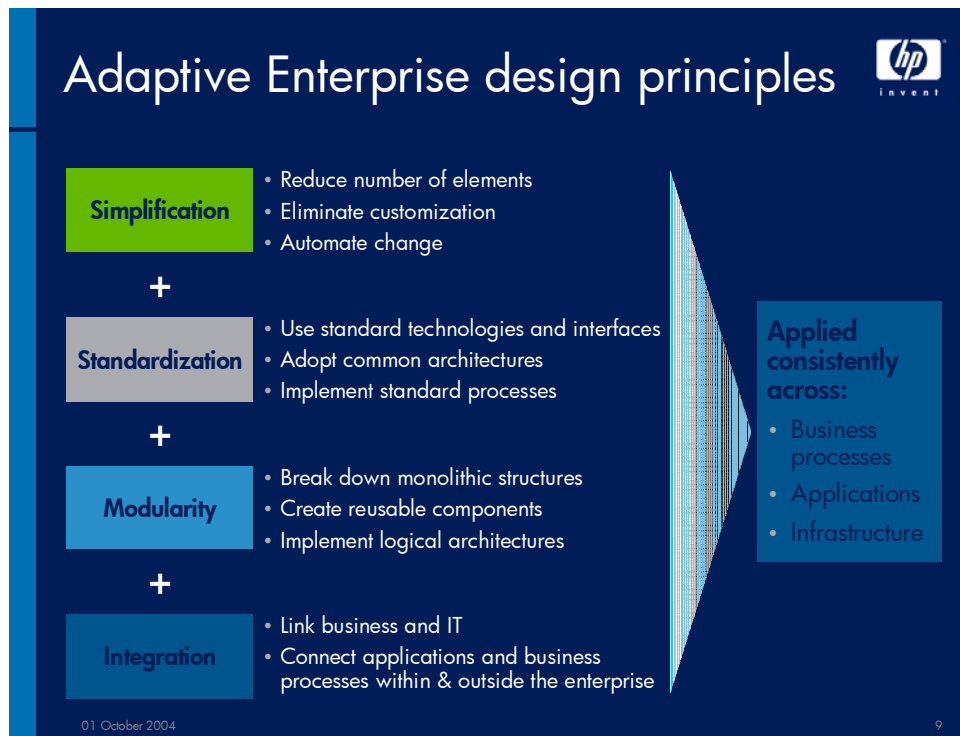


Figure 3 – Adaptive Enterprise design principles

OpenVMS already facilitates -

simplification by allowing easy consolidation and central points of management

standardization with the port to the Itanium Processor Family in common with HP-UX, Linux, Non-Stop and Windows®- 64;

modularity by virtualization of flexible compute resources and

integration with existing/strategic, new and external applications through the use of new Web Services

Adaptive Enterprise Components

The Adaptive Enterprise consists of a number of components, as follows:

Virtualization

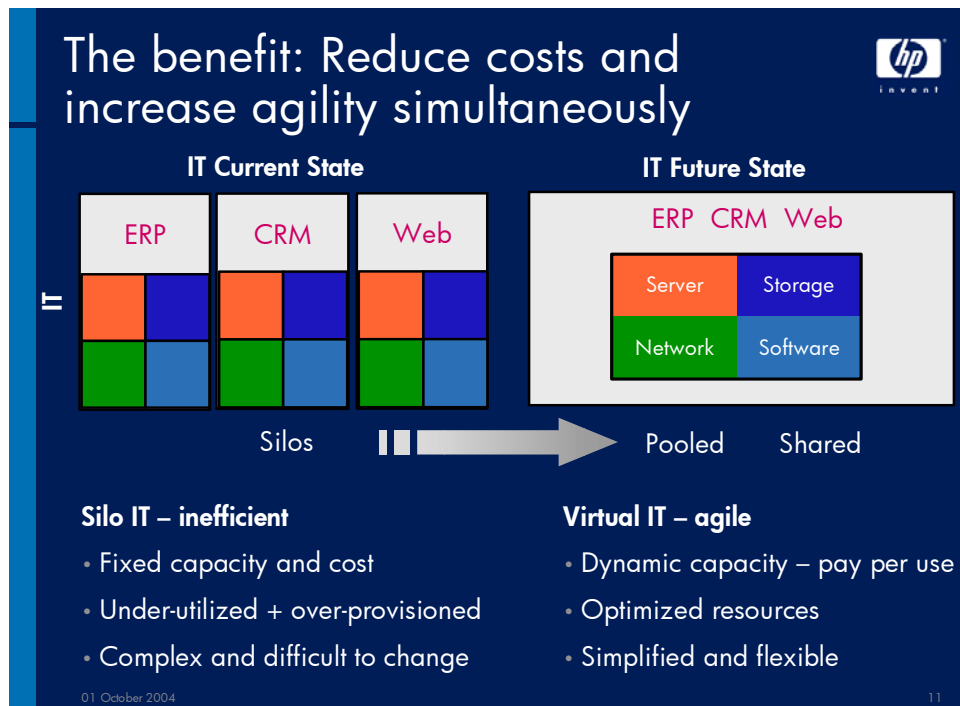


Figure 4 - Instead of discrete systems dedicated to specific applications as is the norm today, business processes have access to a virtual resource pool of services, applications, servers, storage and network

OpenVMS adopted the right hand picture a number of years ago with the deployment of OpenVMS clusters and OpenVMS Galaxy – see later.

Applications

Increasingly, new applications are being written in portable, standardized languages and environments such as Java, XML and J2EE and delivered as services, as opposed to the monolithic applications of the present. These new Web Service tools are also being used to expose long established applications written in e.g. COBOL and Pascal as a web service. This obviates the need to re-write the application and the new service is in production much quicker. All this leads to increased agility for the company in response to changing business needs and is a fundamental pillar of the Adaptive Enterprise.

OpenVMS Web Services

OpenVMS offers a comprehensive suite of Web Services and applications that deliver the following:

- Extending the ROI of existing applications
- Increasing revenue and reducing costs by developing new revenue generating solutions based on existing applications that have been exposed as services
- Reducing both development costs and response time to new requirements by developing new solutions with Web Services
- Ensuring applications running on OpenVMS can be accessed from non-OpenVMS clients and platforms
- Reducing costs by focusing on a single application infrastructure
- Increased revenue and reduced costs when partners and suppliers adopt the same standardized Web based infrastructure utilising environments such as J2EE and WSDL – Web Services Description Language

Web Service products available on OpenVMS include:

- CSWS – a secure, enhanced Apache HTTP web server
- CSWS_JAVA – a.k.a. Tomcat, a Java Server Pages and Servlet engine
- SOAP, being replaced by AXIS – Web services development and deployment toolkits
- Xerces – XML parser available in Java and C++
- Xalan – Transforms XML to HTML and other XML documents
- NetBeans – Integrated Java development environment
- Enterprise Directory – a high performance, secure, distributed directory
- Ant – build environment supplied as part of Tomcat
- BEA Web Logic Server – Enterprise grade J2EE application server
- InterSystems Ensemble – rapid integration and composite application development environment


Applications on established OpenVMS systems are prime candidates for integration by web services. In addition, OpenVMS systems possess valuable business data and “corporate memory” that can now be better accessed and manipulated in response to new business demands.

Infrastructure Services

In the past, resources were delivered by assigning servers directly to applications. But low utilization and high cost from “one application, one server” policies compels organizations toward shared, virtualized, and on-demand solutions that make better use of processor and storage capacity. The HP adaptive infrastructure technologies and solutions accept requirements from and deliver infrastructure services to applications through open industry interfaces and Web services.

For mission-critical enterprise applications, HP has long worked closely with ISVs to develop reference architectures in which application demand patterns are closely matched to specific processor and storage configurations. These architectures, like all communications within the reference architecture, will now be developed using open standards. This commitment ensures that the HP architecture will be the most inclusive environment for multi-vendor hardware, applications and middleware.

Adaptive Enterprise Areas of Focus

HP offerings for an Adaptive Enterprise 	
Areas of focus	BCS and partners
Continuous & secure operations	<ul style="list-style-type: none"> • Servers • Operating systems • High availability
Automated & intelligent management	<ul style="list-style-type: none"> • Clustering • Manageability
Dynamic resource optimization	<ul style="list-style-type: none"> • HP Utility Pricing Solutions • Virtual Server Environment (VSE)
Foundation for the future	<ul style="list-style-type: none"> • Itanium architecture • Business systems evolution

03 November 2004 12

Figure 5 - Areas of focus and HP Business Critical Systems' offerings

Continuous and Secure Operations

In 2002, HP OpenVMS celebrated 25 years of shipments and in this time it has led the industry with respect to security, performance and high availability.




Figure 6 - The OpenVMS Security Model

Security

OpenVMS was designed from the outset with security built-in, as the above diagram shows. As competing operating systems have found, it is not possible to build effective security on top of an established operating system as the Department of Defense, European Defence Ministries and other government agencies will testify. OpenVMS has the lowest CERT Advisory score of any general purpose operating system available today and is two orders of magnitude more secure than some operating systems originating in the 80s and 90s.

Performance

Developed on the world beating DEC VAX architecture, OpenVMS was at the forefront of mini-computer performance in the 80s and early 90s. As compute demands inexorably rose, it was necessary to move to a 64-bit platform which was developed by Digital Equipment Corporation as the Alpha chip, the world's first commercially available true 64-bit design. HP AlphaServer systems regained 64-bit performance leadership with the 64 processor AlphaServer GS1280 server at the beginning of the 21st Century, when the decision was made to move to the Intel® Itanium® Processor Family architecture that was jointly designed with HP. This leads to tangible benefits for existing and new customers on a number of levels – see later.

Clustering, High Availability and Manageability

OpenVMS introduced the world to clustering in 1983 and has since maintained the lead in this technology. OpenVMS clusters enable system managers to configure up to 99.999% uptime by having access to multiple CPUs (up to 96 nodes) and servers sharing replicated storage. By deploying across two or more split-site clusters, the service can continue un-interrupted even if a disaster befalls one of the data-centres.

An excellent report by analysts *Illuminata* in August 2002 declared OpenVMS the “gold standard” for clustering and disaster tolerance compared to leading UNIX implementations from IBM and Sun. The report can be found here: <http://h71000.www7.hp.com/openvms/whitepapers/Illuminata.pdf>

OpenVMS has a comprehensive suite of management tools and is supported by a number of leading management Independent Software Vendors ensuring that all elements of OpenVMS, OpenVMS clusters and the network (DECnet and TCP/IP) are securely managed.

OpenView

In addition, OpenVMS shipped the OpenView Operations Agent V1.0 in the first quarter of 2004. This aids the integration of OpenVMS systems within an HP-UX or Windows Server environment and will be followed by other OpenView components in due course.

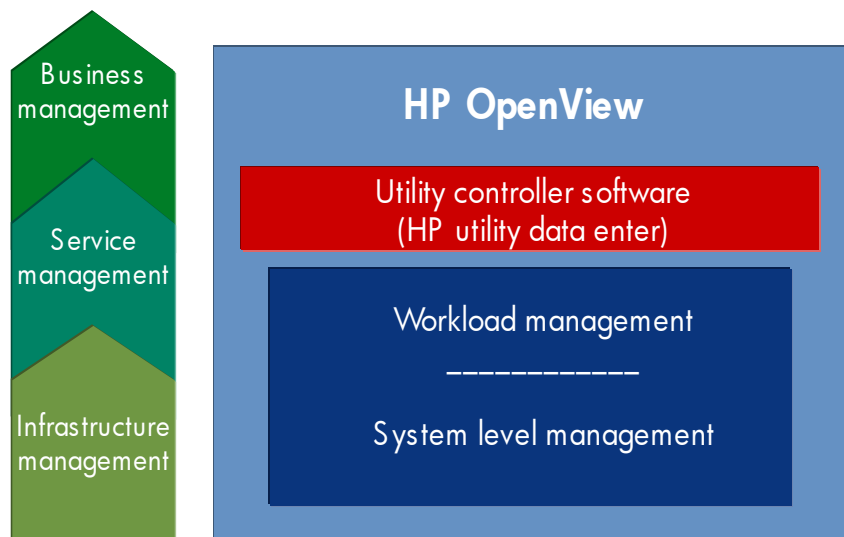


Figure 7 - The role of OpenView in Adaptive Management

Dynamic Resource Optimization

This is the real-time delivery of computing resources when and where needed achieved by soft partitioning. It offers the flexibility to get the best use from resources, investments and systems, today and tomorrow.

Virtual Server Environment - OpenVMS Galaxy

Introduced in 1997, OpenVMS Galaxy soft partitioning has led the way with dynamic CPU resource allocation and optimization enabling the following:

- Responding to Changing Business Requirements
 - Handle multiple workloads
 - Scale to use full computing power
 - Dynamically reallocate CPU power to where it's needed
- Reduce Cost of Ownership through:
 - Server, Application and Management consolidation
 - CPU reassignment
 - Optimization

Multi-CPU systems share resources such as memory and storage at very high speed enabling near instantaneous reconfiguring to satisfy the highest priority jobs and Service Level Agreements (SLAs). An actual OpenVMS disaster tolerant cluster and OpenVMS Galaxy software implementation is shown in the following figure:

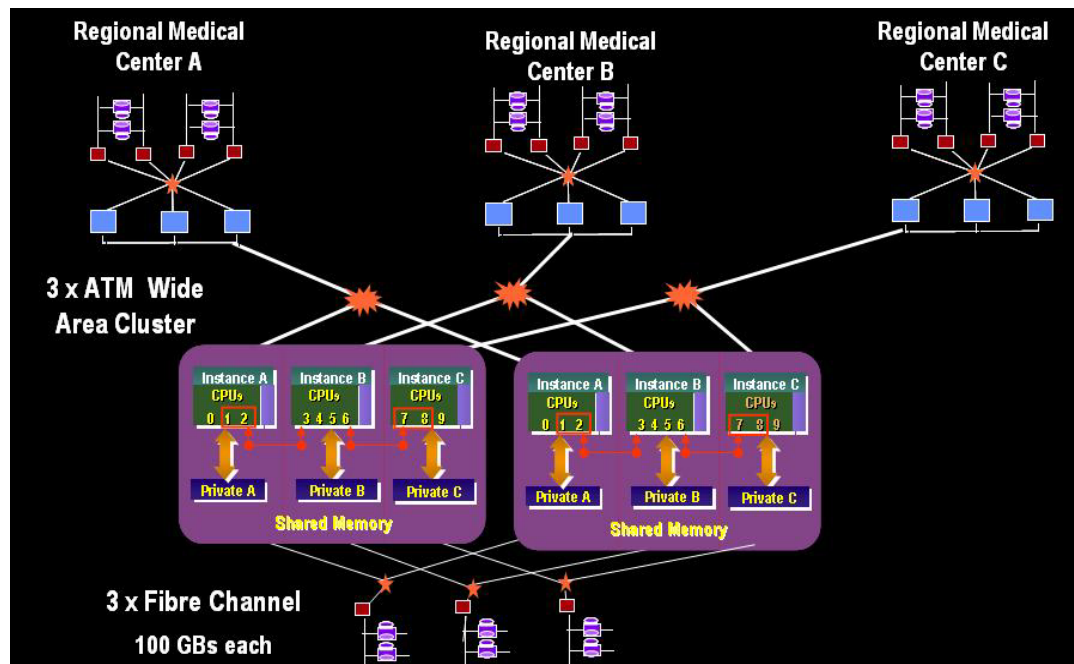


Figure 8 - A comprehensive disaster tolerant deployment of a replicated Central Medical Records Database deployment, using standard OpenVMS systems, OpenVMS cluster software and OpenVMS Galaxy software

Workload Manager

Translating the demands of the Service Level Objectives (SLOs) - themselves generated from the fluctuating process streams of multiple Business Units - is the Workload Manager (WLM). WLM ensures that resources are optimized to satisfy the most pressing SLO requirements first, interacting with the lower level Process Resource Manager (PRM).

Global Workload Manager

A heterogeneous development of WLM has been initiated by HP that will bring the capabilities of WLM to multiple operating systems sharing a single Integrity server or clustered networks of Integrity servers. This product is known as global Workload Manager – gWLM.

OpenVMS Engineering has made detailed design studies of gWLM and is committed to offering gWLM on OpenVMS in the future.

Process Resource Manager – OpenVMS Class Scheduler

The Process Resource Manager – PRM – is the mechanism that guarantees processing time to business critical applications and users based upon their priority derived from the SLOs. Applications may be placed into classes, easing manageability and each class allocated a priority and a percentage of the available CPUs, memory and I/O bandwidth. This mechanism behaves similarly to the *Class Scheduler* that already ships within the OpenVMS operating system.

HP Utility Pricing Solutions

HP has offered a number of purchasing options in combination with HP Finance, the second largest technology finance company in the world. Pay Per Forecast (PPF) is a contractual agreement between HP Finance and the customer that require no engineering within the operating system to implement.

Instant Capacity – iCAP – enables the customer to order active CPUs and iCAP CPUs within a configuration and only get charged full price for the active CPUs. On delivery, email is configured on the system enabling access to the HP iCAP Billing System. When the customer wishes to activate some or all of the the iCAP CPUs, upon receipt of an email from the customer HP sends an Authorization Key allowing the customer to turn on the defined number of iCAP CPUs. At this point the customer is billed for the newly activated CPUs.

A variant of iCAP known as Temporary iCAP – TiCAP – allows the customer to activate iCAP CPUs for up to 30 days for testing, certification etc., purposes.

OpenVMS AlphaServers have offered Instant Capacity and Temporary Instant Capacity for a number of years and they are available on the ES80 and GS1280 AlphaServer systems.

Pay Per Use (PPU) offers detailed measurement of CPU utilisation throughout the 24 hours, ensuring that customers only pay for what they use. PPU comes in two flavours:

- **Active CPU** – CPU status (active or turned off) is monitored and recorded every five minutes and the customer is charged only for the average of the active CPUs over the 24 hours. This variant of PPU is favoured by ISVs that charge for their software based on the number of CPUs on which it runs

- **Percent CPU** - All CPUs are running at all times to ensure good response times and customers are only charged for the collective percentage utilisation of active CPUs throughout the 24 hours. This variant of PPU is favoured by customers as all CPUs are utilized at all times and there is no danger of inadvertently turning off CPUs which may compromise applications

OpenVMS Engineering is conducting detailed design studies of both implementations and OpenVMS is committed to delivering both variants in the future.

Per Processor Licensing

In addition, OpenVMS is actively implementing Per Processor Licensing (PPL) on Integrity servers, whereby customers are charged for OpenVMS based on the number of processors they have installed in their systems. Therefore, a customer purchasing an eight CPU Superdome with racks to grow to 64 CPUs and requiring the full complement of processors in the future will only be charged initially for OpenVMS on the eight CPUs. PPL will debut in 8.2 at end 2004.

Foundation for the Future

OpenVMS has a very healthy future now that it is coupled to the Intel® Itanium family of 64-bit chips. This is how OpenVMS fits in with the overall HP operating system family on the Intel® Itanium®- based HP Integrity server family

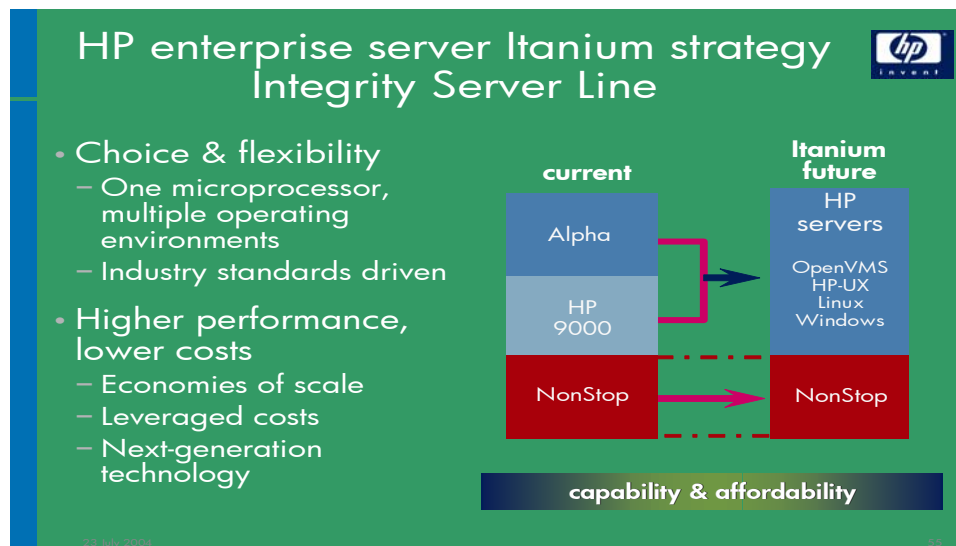


Figure 9 - OpenVMS as an integral member of the HP operating system family on the Intel® Itanium® processor

Moving all HP operating systems to the HP Integrity server family offers economy of scale at the system and chip level, financially benefiting HP's customers. The operational benefits of sharing common CPUs and HP Integrity servers, has been outlined above.

The OpenVMS Adaptive Enterprise Timeline

2004: OpenVMS can deliver today on many aspects of the Adaptive Enterprise:

- Support of industry standard web services such as J2EE, Secure Web Server (enhanced Apache), Tomcat, WSDL, JNDI, LDAPv3, XML, SOAP etc.
- Highest levels of security in the industry
- Industry leading Disaster Tolerance, High Availability and Clustering
- Network-wide Authentication
- Virtualization
- System and CPU utilization optimization with OpenVMS Galaxy
- Comprehensive Management tools
- Flexible pricing with standard leasing and Pay Per Forecast

1H2004

- In the first Quarter of 2004 OpenVMS delivered an OpenView Operations Agent for OpenVMS V1.0. This enables real-time OpenVMS AlphaServer system monitoring and event management utilizing native agents and Smart Plug-Ins (SPIs) and also monitors the performance, detects bottlenecks and reports on the availability of the OpenVMS systems and can send alarms and notifications

2H2004

- With the first release of production OpenVMS 8.2, OpenVMS will ship on the HP Integrity Server Line system family and coexist with HP-UX, Linux and Windows64 on the same platform.
- Per Processor Licensing will also be available with OpenVMS 8.2.

2005 and beyond

- Global Workload Manager (gWLM) support using a common code base with HP-UX and Linux, enabling customers to satisfy multiple SLAs and maximize the utilization of multiple Integrity systems and clusters from a common interface
- Pay Per Use – Percent CPU and Active CPU support – ensuring customers only pay for the compute resources they consume
- Continued investigations into other Virtualization technologies such as Virtual Machine (VM) and Firm Partitioning

HP VSE for HP OpenVMS Roadmap

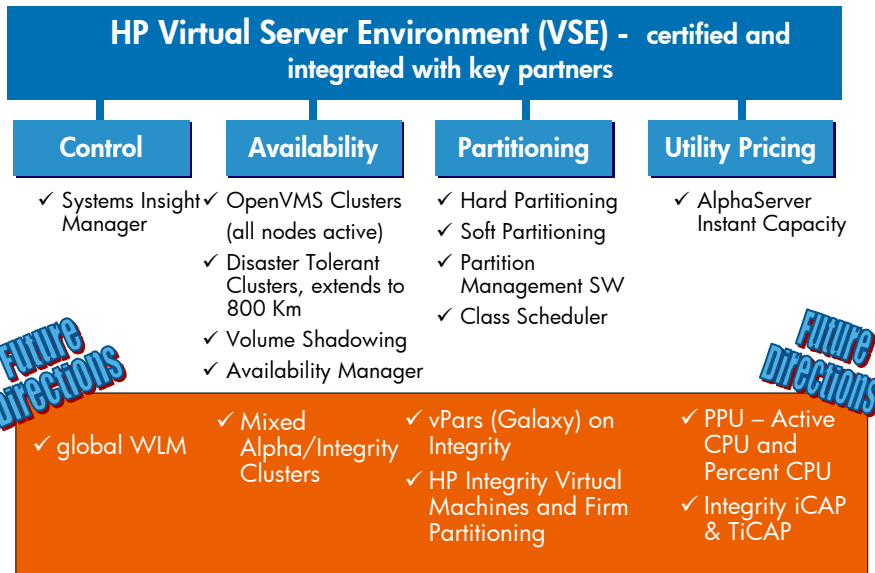


Figure 10 – OpenVMS capabilities enabling the Virtual Server Environment

Adaptive Enterprise Components		1H04	2H04	1H05	2H05	1H06	2H06
Web Services	XML, SOAP, WSDL, J2EE etc.						
Adaptive Management	OpenView and SPIs Multiple 3rd Parties						
HP Utility Pricing Solutions	Pay Per Forecast AlphaServer Instant Capacity Integrity Instant Capacity (iCAP) Pay Per Use (PPU) Per Processor Licensing (PPL)						
Virtualization Technology	Soft Partitions vPars (Galaxy) Hard Partitions nPars Global Workload Manager (gWLM) Process Resource Mgr (PRM) Virtual Machine						
Continuous Secure Operations	High Availability N-site Disaster Tolerance & Recovery Active-Active Clusters Network-wide Authentication C2 Cluster-wide domains						
Platform	Itanium Alpha						

Indicates the phasing in of this component

Figure 11 - The OpenVMS Adaptive Enterprise Components and Timeline

Conclusion

OpenVMS has a very bright future within HP because of its strong affinity with the Adaptive Enterprise vision. As described above, OpenVMS already delivers on many of the Adaptive Enterprise components and with its comprehensive support of Web Services, ensures that existing applications and data can be utilized better and faster in response to changing business demands.

Plans for OpenVMS include development of specific Adaptive Enterprise components such as OpenView, Smart Plug Ins, Global Workload Manager and Pay Per Use ensuring that, over time, OpenVMS is a key pillar of the Adaptive Enterprise portfolio.

Fundamental attributes of OpenVMS such as high levels of security; high availability and disaster tolerance through OpenVMS clustering; adaptability and agility through soft partitioning (OpenVMS Galaxy); strong performance on both AlphaServer systems and the Integrity server family, will never go out of fashion. HP is committed to building upon this peerless foundation to ensure OpenVMS better serves existing and new customers as part of the HP Adaptive Enterprise.

For more information

<http://h71000.www7.hp.com/>

HP Partitioning Continuum, HP, 2002

HP ENSAextended technical overview, HP, 2003

Call to action

HP ENSAextended technical overview, HP, 2003

© 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Itanium is a trademark or registered trademark of Intel Corporation in the U.S. and other countries and is used under license.

11/2004